MYCIN: A KNOWLEDGE-BASED COMPUTER PROGRAM APPLIED TO INFECTIOUS DISEASES*

Edward H. Shortliffe
Department of Medicine
Stanford University School of Medicine
Stanford, California 94305

A rule-based expert system is described which uses artificial intelligence techniques, and a model of the interaction between physicians and human consultants, to attempt to satisfy the demands of a user community that is often reluctant to experiment with computer technology. Experience to date has demonstrated that the program is efficient, relatively easy to use, and reliable in the domain of bacteremia therapy selection. Future work will involve broadening and evaluating the program's expertise in other areas of infectious disease therapy. To that end rules regarding diagnosis and treatment of meningitis have been written and are currently under evaluation.

Introduction

Few potential user populations are as demanding of computer technology as are practicing physicians. This is due to a variety of factors which include the physician's independence as a lone decision maker, the seriousness with which he views actions that may often have life-and-death significance, and the overwhelming time demands which tend to make him impatient with any innovation that breaks up the finely-tuned flow of his daily routine. Yet as medical science has expanded, the individual practitioner has become increasingly less able to manage all the expertise he needs if he is to provide modern medical care. Consultation from subspecialists has therefore become a common and accepted part of practice for those physicians fortunate enough to have easy access to the kinds of expertise they need. Away from large urban or academic centers such consultative advice may be more difficult to obtain, and it is partly for this reason that efforts have been made to develop computer programs with sufficient subspecialty expertise to function in a reliable consultative role.

Despite the medical professions' common reluctance to experiment with clinical computing[1], many consultation programs have failed to emphasize the development of mechanisms for encouraging their use by physicians. In designing the MYCIN program, a clinical consultation system, we have attempted to recognize the need to place the ultimate medical decisions in the hands of the physician. MYCIN includes mechanisms so that the doctor may understand not only the program's advice, but also the basis on which the relevant clinical decisions were reached. Such an understanding is encouraged by allowing the physician to maintain the initiative throughout the consultation, with an ability to request clarification or justification of puzzling points along the way. This process parallels the familiar form of dialog and advice from a human consultant (Fig. 1) and is thereby less threatening than a system which simply submits dogma. MYCIN's area of expertise is the selection of antimicrobial therapy for patients with severe infections. This paper briefly summarizes the MYCIN system and discusses its debt to symbolic reasoning techniques from the field of artificial intelligence.

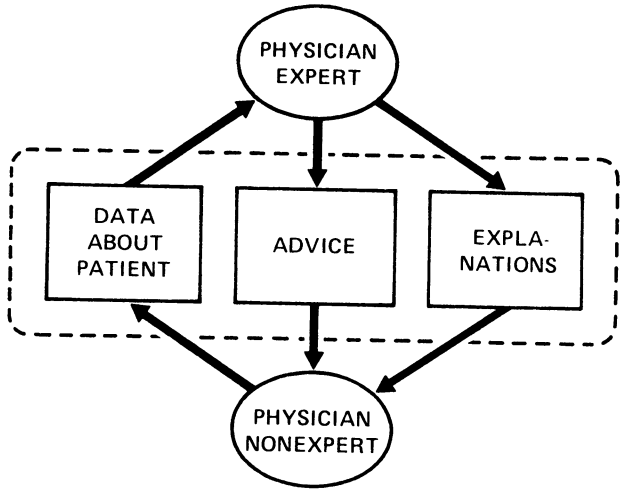

Figure 1 - Diagram summarizing the flow of information between physician and expert in the human consultation process. (Figure reproduced from reference 10).

Artificial Intelligence

Artificial intelligence (AI) is a subfield of computer science in which the emphasis of the research is on symbolic reasoning rather than numberical computations. In many respects the name "artificial intellingence" is unfortunate because it conjures up threatening images of superhuman machines that challenge those capabilities of mankind that have long been thought to be uniquely human. A recent book by a respected computer expert[2] has examined the field in detail and has generated a great deal of discussion not only about what computers could do but should do[3]. Techniques developed by researchers in AI, however, have great potential for application to problems in medical decision making. Since 1970 a small number of researchers, most of whom have had experience rooted in traditional computational approaches to medical decision making, have begun to recognize the need for computer "understanding" of the clinical domain of interest if useful diagnostic tools are to be developed[4]. They point out that capabilities such as heuristic search through large numbers of possible decisions or actions, computational approaches to natural language understanding and generation, and issues of representation for inferential knowledge (all major AI research areas) are central to the problem of giving a computer program a knowledge of medicine at a conceptual level. Such capabilities might allow the program a spectre of creativity and "humaness" that would in turn make it an acceptable tool for physicians. Several recent research projects have therefore begun to examine the use of symbolic reasoning techniques in medical domains [5-8]. An overview of AI and its relationship to medical problem solving is available elsewhere [9].

---

* This paper was originally prepared for presentation at the Annual Meeting of the Society for Computer Medicine, Las Vegas, Nevada, 10 November 1977.

The MYCIN System

MYCIN is a LISP program designed to serve as a clinical consultant on the subject of therapy selection for patients with infections. The program may be envisioned as interposed between the expert and nonexpert in much the way that the large box is positioned in Fig. 1. The difference is that the human expert can offer only general knowledge to the program, not patient-specific decisions. The program thus becomes the decision maker, using general medical knowledge from experts to assess a specific patient and to give advice plus explanations for its judgments.

Fig. 2 details the organization of MYCIN relative to the human consultation process depicted in Fig. 1. As before, the nonexpert offers data about his patient and in return receives both advice and, when desired, information via one of two internal explanation mechanisms (the "General Question-Answerer" or the "Reasoning Status Checker"). The basis for all decisions is domain-specific knowledge acquired from experts ("Static Knowledge"). A group of computer programs (the "Rule Interpreter") uses this knowledge, and data about the specific patient, to generate conclusions and, in turn, therapeutic advice. It simultaneously keeps a record of what has happened, and this record is available to the explanation routines if the physician asks for justification or clarification of some conclusion the program has reached. The details of this system organization have been discussed in another publication [10].

The three principal design goals which motivated the design shown in Fig. 2 were that MYCIN should a) give clinically useful advice, b) explain decisions when asked to do so, and c) acquire domain-specific knowledge directly from experts. Thus there are three interrelated parts to the MYCIN system [8,11]:

(1) A Consultation System that uses the knowledge base, along with patient-related data entered by the physician to generate therapeutic advice (see references 8-12);

(2) An Explanation System (the general question-answerer and the reasoning status checker) that explains the reasoning used during the consultation and documents the motivation for questions asked or the rationale for conclusions reached (see reference 12); and

(3) A Knowledge Acquisition System that enables experts in antimicrobial therapy to update MYCIN's static knowledge base without requiring that they know how to program a computer (see references 11, 14).

Evaluation

It is important to avoid premature introduction of a program in its real-world environment. This is perhaps particularly true in medicine, where physician resistance to computer-assisted decision making is already so great that an "experimental" system which performs poorly is apt to generate insurmountable biases, regardless of how good its performance eventually becomes. Before MYCIN can be introduced for clinical use, we therefore believe it must be shown that the program recommends either the same therapy as experts do, or recommends an alternate, equally acceptable regimen. This issue breaks down into a number of components which must be analyzed separately: does the program adequately decide whether the patient needs treatment in the first place? does it ask any irrelevant questions? does it ignore important questions? does it correctly conclude what bacteria are apt to be causing the patient's disease? is its therapeutic regimen appropriate for the bacterial pathogens it feels may be present?

In an effort to answer these questions regarding MYCIN's performance in the domain of bacteremia therapy, a prospective study of 15 consecutive patients with positive blood cultures was undertaken. The same set of patient data was provided both to MYCIN and to ten experts in infectious disease therapy. The evaluation data are presented in detail elsewhere [15], but in general they show that the program provides good advice most of the time and that its performance closely approaches that of subspecialists in the clinical domain.

Limitations

One of MYCIN's principal limitations is shared by all clinical consultation programs - - it requires that the physician take the initiative in asking for an interactive session. Despite attention that may be paid to "human engineering" issues during the development of such programs, physicians seldom choose to use computers for tasks they feel they can do themselves[1]. Evidence that the medical profession may be failing to perform a job optimally (reviewed in reference 15 for MYCIN's domain) may not be sufficiently compelling for the individual practitioners to seek out a program's advice. Programs which can monitor physician prescribing habits automatically, and generate warnings when appropriate [16], may thus be more likely to influence physicians and the general quality of patient care.

The MYCIN approach is also limited by its complex control structure and representation scheme. These are probably unnecessary for medical inference tasks in which uncertainty and incomplete knowledge are uncommon occurrences. Until the MYCIN formalism has been tested on other medical decision problems, however, it will be difficult to predict what proportion might be handled more directly by other techniques. Furthermore, the efficiency advantages of approaches relying more heavily on statistical theory must be weighed against the importance of the natural mechanism for explanation and knowledge modification that may be achieved through MYCIN's representation of knowledge in the form of production rules [17].

Another limitation results because MYCIN currently requires more computing power and memory than is reasonable to expect to find in most hospital computing environments. Although the DEC PDP-10 on which we currently operate has been a powerful research machine for development of MYCIN, we are hopeful that the program may eventually be adapted to run on smaller machines. Alternatively, networking technology may allow a single large machine to be shared by a large number of clinical users over a wide geographic area. At present, however, the system functions solely in a research environment, and long range cost estimates are therefore not feasible.

Acknowledgements

The MYCIN Project is supported by BHSRE Grant No. HS01544. Research on the system has been interdisciplinary and has involved the following individuals, all of whom have made substantial contributions to system development and evaluation: J.S. Aikins, B.G. Buchanan, W.J. Clancey, R. Davis, L.M. Fagan, A.C. Scott, and W.J. van Melle (computer scientists), S.G. Axline, S.N. Cohen, F.S. Rhame, and V.L. Yu (physicians), and S.M. Wraith (clinical pharmacist). The research described has been undertaken on the SUMEX computing facility at Stanford University School of Medicine, supported by DRR/NIH Grant RR-00785.

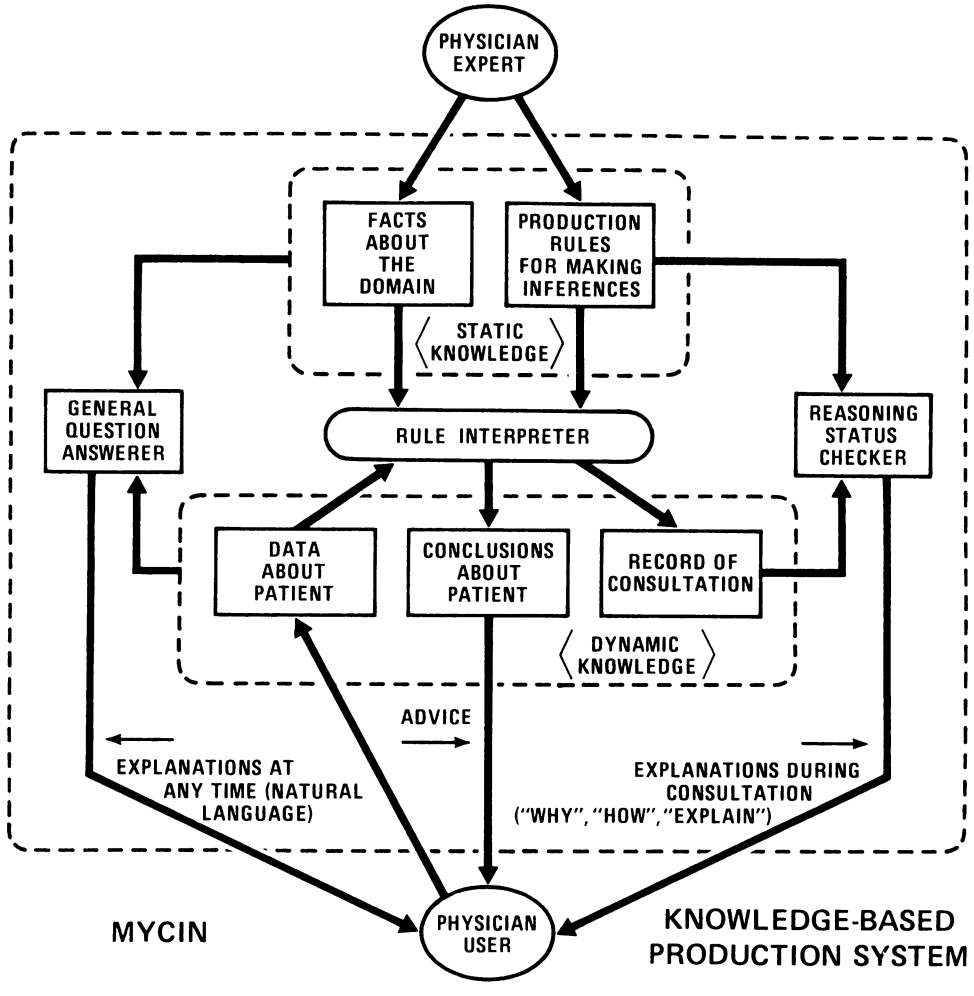

Figure 2 - Diagram summarizing the organization and flow of information within MYCIN. The correlation between this design and the human consultation process depicted in Fig. 1 is discussed in the text. (Figure reproduced from reference 10).

References

1) Startsman, TS and Robinson RE: The attitudes of medical and paramedical personnel towards computers. Comput. Biomed. Res. 5,218-227 (1972).

2) Weizenbaum, J: Computer Power and Human Reason. WH Freeman and Co., New York, 1976.

3) "Experts argue whether computers could reason, and if they should". p. 1 The New York Times, 8 May 1977.

4) Gorry, GA: Computer-assisted clinical decision making. Meth. Inform. Med. 12, 45-51 (1973).

5) Kulikowski, CA, Weiss, S, and Safir A: Glaucoma diagnosis and therapy by computer. Proc. Annu. Meet. Ass. Res. Vision Ophthamol. (May 1973).

6) Pauker SG, Gorry GA, Kassirer JP, and Schwartz WB: Towards the simulation of clinical cognition: taking the present illness by computer. Amer. J. Med. 60, 981-996 (1976).

7) Pople H, Meyers J and Miller R: DIALOG, a model of diagnostic logic for internal medicine. Proc. Fourth Internl. Joint Conf. Artificial Intelligence Tbilisi, USSR, pp. 848-855, September 1975.

8) Shortliffe EH, Davis R, Buchanan BG, Axline SG, Green CC,and Cohen SN: Computer-based consultations in clinical therapeutics-explanation and rule-acquisition capabilities of MYCIN system. Comput. Biomed. Res. 8, 303-320 (1975).

9) Shortliffe EH: Computer-Based Medical Consultations: MYCIN. American Elsevier Publishing Company, Inc., New York, 1976.

10) Shortliffe EH: A rule-based approach to the generation of advice and explanations in clinical medicine. Proc. Workshop Computational Linguistics in Medicine Uppsala, Sweden, May 1977 (in press).

11) Davis R, Buchanan BG and Shortliffe EH: Production rules as a representation for a knowledge-based consultation program. Artificial Intelligence 8, 15-45 (1977)

12) Shortliffe EH and Buchanan BG: A model of inexact reasoning in medicine. Math. Biosci. 23; 351-379 (1975).

13) Scott AC, Clancey WJ, Davis R and Shortliffe EH: Explanation capabilities of knowledge-based production systems. Amer. J. Computational Linguistics, Microfiche 62, 1977. Also available as TR HPP-77-1, Heuristic Programming Project, Stanford University,March 1977.

14) Davis R: Applications of Meta Level Knowledge to the Construction, Maintenance, and Use of Large Knowledge Bases. Doctoral dissertation, Stanford University, June 1976. Also available as AIM-283 Stanford Artificial Intelligence Laboratory, Stanford University, July 1976.

15) Buchanan BG, Davis R, Yu, VL, and Cohen SN: Rule-based medical decision making by computer. Proceedings of MEDINFO 77, Toronto, Canada, August 1977 (in press).

16) Cohen SN, Armstrong MF, Briggs RL et al: Computer-based monitoring and reporting of drug interactions. Proceedings of MEDINFO 74, pp. 889-894, Stockholm, Sweden, August 1974.

17) Davis R and King J: An overview of production systems. In Machine Representations of Knowledge (Machine Intelligence 8, eds. E.W. Elcock and D. Michie), John Wylie, New York, 1976. Also available as AIM-271, Stanford Artificial Intelligence Laboratory, Stanford University, October 1975.